arXiv:cs/0507020v1 [cs.LO] 7 Jul 2005

First-order queries on structures of bounded degree are computable with constant delay

Arnaud Durand * Etienne Grandjean †

February 1, 2008

Abstract

A bounded degree structure is either a relational structure all of whose relations are of bounded degree or a functional structure involving bijective functions only. In this paper, we revisit the complexity of the evaluation problem of not necessarily Boolean first-order queries over structures of bounded degree. Query evaluation is considered here as a dynamical process. We prove that any query on bounded degree structures is CONSTANT-DELAY$_{lin}$, i.e., can be computed by an algorithm that has two separate parts: it has a precomputation step of linear time in the size of the structure and then, it outputs all tuples one by one with a constant (i.e. depending on the size of the formula only) delay between each. Seen as a global process, this implies that queries on bounded structures can be evaluated in total time $O(f(|\varphi|).(|\mathcal{S}| + |\varphi(\mathcal{S})|))$ and space $O(f(|\varphi|).|\mathcal{S}|)$ where $\mathcal{S}$ is the structure, φ is the formula, $\varphi(\mathcal{S})$ is the result of the query and f is some function.

Among other things, our results generalize a result of [See96] on the data complexity of the model-checking problem for bounded degree structures. Besides, the originality of our approach compared to that [See96] and comparable results is that it does not rely on the Hanf's model-theoretic technic (see [Han65]) and is completely effective.

Introduction

Evaluating the expressive power of logical formalisms is an important task in theoretical computer science. It has many applications in numerous fields such as complexity theory, verification or databases. In this latter case, it often amounts to determine how difficult it is to compute a query written in a given language. In this vein, determining which fragments of first-order logic defines tractable query languages has deserved much attention.

It is well known, that over an arbitrary signature, computing a first-order query can be done in time polynomial in the size of the structure (and even in logarithmic space and AC^0). However the exponent of this polynomial depends heavily on the formula

*LACL - CNRS FRE 2673, Département d'informatique, Université Paris 12, 94010 Créteil - France. Email: durand@univ-paris12.fr

†GREYC - CNRS UMR 6072, Université de Caen - Campus 2, F-14032 Caen cedex - France. Email: grandjean@info.unicaen.fr

size (more precisely, on the number of variables). Nevertheless, for particular kinds of structures or formulas the complexity bound can be substantially improved. In [See96], it is proved that checking if a given first-order sentence φ is true (i.e., the Boolean query or model-checking problem) in a structure $\mathcal{S}$ all of whose relations are of bounded degree can be done in linear time in the size of $\mathcal{S}$. The method used to prove this result relies on old model-theoretic technics (see [Han65]). It is perfectly constructive but hardly implementable. Later, still using such kind of methods, several other tractability results have been shown for the complexity of the model-checking of first-order formulas over structures or formulas that admit nice (tree) decomposition properties (see [FFG02]).
In this paper, a *bounded degree structure* is either a relational structure all of whose relations are of bounded degree or a functional structure involving bijective functions only.
The main goal of this paper is to revisit the complexity of the evaluation problem of not necessarily Boolean first-order queries over structures of bounded degree. We regard query evaluation as a *dynamical process*. Instead of considering the cost of the evaluation globally, we measure the delay between consecutive tuples, i.e., query problems are viewed as enumeration problems. This latter kind of problems appears widely in many areas of computer science (see for example [EG95, EGM03, BGKM00, KSS00, Gol94] or [JYP88] for basic complexity notions on enumeration). However, to our knowledge, relation to query evaluation has not been investigated so far.
We prove that any query on bounded degree structures is CONSTANT-DELAY$_{lin}$, i.e., can be computed by an algorithm that has two separate parts: it has a precomputation step whose time complexity is linear in the size of the structure and then, outputs all the solution tuples one by one with a constant (i.e., depending on the size of the formula only) delay between two successive tuples. Seen as a global process, this implies that queries on bounded structures can be evaluated in total time $O(f(|\varphi|).(|\mathcal{S}|+|\varphi(\mathcal{S})|))$ and space $O(f(|\varphi|).|\mathcal{S}|)$ where $|\mathcal{S}|$ is the size of the structure $\mathcal{S}$, $|\varphi|$ is that of the formula φ, $|\varphi(\mathcal{S})|$ is the size of the result $\varphi(\mathcal{S})$ of the query and f is some function. As a corollary, it implies that the time complexity of the model-checking problem is $O(f(|\varphi|).|\mathcal{S}|)$ thus providing an alternative proof of the result of [See96].
A particularity of the main method used in this paper is that it does not rely on model-theoretic technic as previous results of the same kind (see, for example, [See96] or [Lin04] for a generalization to least-fixed point formulas). Instead, we develop a quantifier elimination method suitable for bijective unary functions and apply it to obtain our complexity bound. An advantage of this method is that it is effective and easily implementable. Another advantage is that our paper is completely self-contained.
Besides, the CONSTANT-DELAY$_{lin}$ class is an interesting notion by itself and is, to our knowledge, a new complexity class for enumeration problems: as proved for linear time complexity (the class DLIN studied in [GS02]) it can be shown that CONSTANT-DELAY$_{lin}$ is a robust class and is in some sense the minimal robust complexity class of enumeration problems.
The paper is organized as follows. First, basic definitions are given in Section 1. In particular, in Subsection 1.3, we recall definitions about enumeration problems and introduce the notion of constant delay computation and prove some basic properties about it. In Section 2, the quantifier elimination method is introduced and is applied to the evaluation problem of first-order formulas over functional structures all of whose func-

tions are bijective. In Section 3, using classical logical interpretation technics, this later problem is reduced in linear time to the first-order query problem over structures of bounded degree thus providing the same bound for it. Finally, in Subsection 3.3, consequences about the complexity of the subgraph (resp. induced subgraph) isomorphism problem are given.

1 Definitions

1.1 Logical definitions and query problems

We suppose the reader to be familiar with basic notions of first-order logic. A *signature* σ is a finite set of relational and functional symbols of given arities (0-ary function symbols are constants symbols). The arity of σ is the maximal arity of its symbols. The set σ is called *unary functional* if all its symbols are of arity bounded by one.
A (finite) σ-structure consists of a domain D together with an interpretation of each symbol of σ over D (the same notation is used here for each signature symbol and its interpretation).
In this paper, we will distinguish between two kinds of signatures on which semantical restrictions on their possible interpretation are imposed:

- Either σ is made of constant and monadic (i.e., unary) relation symbols and unary function symbols whose interpretation is taken among bijective functions (i.e., permutations) only,
- Or σ contains relation symbols only whose degrees are bounded by some given constant (detailed definitions about bounded degree relations are delayed till section 3).

Structures defined by either of semantical restrictions will be called *bounded degree structures*.
In what follows we make precise notions and problems about first-order logic over bijective structures.

Definition 1 *Let $\sigma = \{\overline{c}, \overline{U}, f_1, \ldots, f_k\}$ be a signature consisting of constant symbols $c_i \in \overline{c}$, of monadic predicates $U_i \in \overline{U}$ and of unary function symbols f_i, $i = 1, \ldots, k$. A* bijective σ-structure *is a σ-structure $\mathcal{S}$ of the form $\mathcal{S} = \langle D; \overline{c}, \overline{U}, f_1, \ldots, f_k \rangle$ where each f_i is a permutation on domain D.*

One of the main results of this paper provides a quantifier elimination method over bijective structures. As it is usual for such kind of result, the elimination will be done in a richer language. The following definition is required.

Definition 2 *A* bijective term $\tau(x)$ *is of the form $f_1^{\epsilon_1} \ldots f_l^{\epsilon_l}(x)$ where $l \geq 0$, x is a variable and where each $f_i^{\epsilon_i}$ is either the function symbol f_i or its reciprocal f_i^{-1}. The term $\tau^{-1}(x)$ denotes the reciprocal of the term $\tau(x)$.*
A bijective atomic formula *is of one of the following four forms where $\tau(x)$ and $\tau_1(x)$ are bijective terms:*

- *either a bijective equality* $\tau(x) = \tau_1(y)$*,*
- *or* $\tau(x) = c$ *where* c *is a constant symbol,*
- *or* $U(\tau(x))$ *where* U *is a monadic predicate,*
- *or a cardinality statement* $\exists^k_x \Psi(x)$ *where the quantifier* $\exists^k_x$ *is interpreted as "there exist at least* k *values of* x *such that" and* Ψ *is a Boolean combination of bijective atoms* $\alpha(x)$ *over variable* x *only.*

As the reciprocal of each function symbol can be used, each bijective equality $\tau(x) = \tau_1(y)$ can be rephrased as $\tau_2(x) = y$ where $\tau_2(x) = \tau_1^{-1}\tau(x)$. A *bijective literal* is a bijective atomic formula or its negation.

Definition 3 *The set* $\mathbf{FO_{Bij}}$ *of bijective first-order formulas is the set of first-order formulas built over bijective atomic formulas of some unary signature* σ*.*

Let $\overline{t} = (t_1, \ldots, t_k)$ be a k-tuple of variables and $\varphi(\overline{t})$ and $\varphi'(\overline{t})$ be two σ-formulas with free variables $\overline{t}$. Formulas $\varphi(\overline{t})$ and $\varphi'(\overline{t})$ are *equivalent* if for all σ-structures $\mathcal{S}$ and all tuples $\overline{a}$ of element of the domain with $|\overline{a}| = |\overline{t}|$ it holds that:

$$(\mathcal{S}, \overline{a}) \models \varphi(\overline{t}) \text{ iff } (\mathcal{S}, \overline{a}) \models \varphi'(\overline{t}).$$

In this paper query problems are considered for specific classes of first-order formulas (and structures). One of the specific problems under consideration here is the following.

QUERY($\mathbf{FO_{Bij}}$)
Input: a unary functional signature σ, a bijective σ-structure $\mathcal{S}$ and a first-order bijective σ-formula $\varphi(\overline{x})$ with k free variables $\overline{x} = (x_1, \ldots, x_k)$
Parameter: φ
Output: $\varphi(\mathcal{S}) = \{\overline{a} \in D^k : (\mathcal{S}, \overline{a}) \models \varphi(\overline{x})\}$.

The Boolean query problem (the subproblem where $k = 0$) is often called a model-checking problem. It will be denoted by MC($\mathbf{FO_{Bij}}$) here. As suggested by the formulation of the query problem, we are interested in its parameterized complexity and the complexity results given here consider the size of the query formula φ as the parameter (see [DF99]).

1.2 Model of computation and measure of time

The model of computation used in this paper is the Random Access Machine (RAM) with uniform cost measure (see [AHU74, GS02, GO04, FFG02]). As query problems are the main subject of this paper, instances of problems always consist of two kinds of objects: first-order structures and first-order formulas.
The *size* $|I|$ of an object I is the number of registers used to store I in the RAM. If E is the set $[n]$, $|E| = card(E) = n$. If $R \subseteq D^k$ is a k-ary relation over domain D, with $|D| = card(D)$, then $|R| = k.card(R)$: all the tuples $(x_1, \ldots, x_k)$ for which $R(x_1, \ldots, x_k)$ holds must be stored, each in a separate k-tuple of registers. Similarly, if f is a unary function from D to D, all values $f(x)$ must be stored and $|f| = |D|$.

If φ is a first-order formula, $|\varphi|$ is the number of occurrences of variables, relation or function symbols and syntactic symbols: $\exists, \forall, \wedge, \vee, \neg, =, "(", ")", ",", "$. For example, if $\varphi \equiv \exists x \exists y\ R(x,y) \wedge \neg(x = y)$ then $|\varphi| = 17$.

All the problems we consider in this paper are parameterized problems: they take as input a list of objects made of a σ-structure $\mathcal{S}$ and a formula φ and as output the result of the query size $\varphi(\mathcal{S})$. Due to the much larger size, in practice, of the structure $\mathcal{S}$ than the size of formula φ, $|\mathcal{S}| >> |\varphi|$, this latter one, $|\varphi|$, in considered here as the parameter.
A problem **P** is said to be computable in time $f(|\varphi|).T(|\mathcal{S}|, |\varphi(\mathcal{S})|)$ for some function $f : N \rightarrow R^+$ if there exists a RAM that computes **P** in time (i.e., the number of instructions performed) bounded by $f(|\varphi|).T(|\mathcal{S}|, |\varphi(\mathcal{S})|)$ using space, i.e., addresses and register contents also bounded by $f(|\varphi|).T(|\mathcal{S}|, |\varphi(\mathcal{S})|)$. The notation $O_\varphi(T(|\mathcal{S}|, |\varphi(\mathcal{S})|))$ is used when one does not want to make precise the value of function f. It is also assumed that the function T is at least linear and at most polynomial, i.e., $T(n,p) = \Omega(n+p)$ and $T(n,p) = (n+p)^{O(1)}$. To give an example and to relate our complexity measure to the logarithmic cost measure, in case T is linear, i.e., $T(n,p) = n+p$, the number of bits manipulated by the RAM is well linear in the number of bits needed to encode the input and the output.

1.3 Enumeration algorithms and constant delay computation

In this section, A is a binary predicate. Enumeration problems will be defined by reference to such a predicate.

Definition 4 *Given a binary relation A, the enumeration function* ENUM$\cdot A$ *associated to A is defined as follows. For each input x:*

$$\text{ENUM}\cdot A(x) = \{y \ :\ A(x,y) \textit{ holds } \}$$

Remark 1 *Query problems may evidently be seen as enumeration problems. The input x is made of the structure $\mathcal{S}$ and the formula $\varphi(\overline{x})$, a witness y is a tuple $\overline{a}$ and evaluating predicate A amounts to check whether $(\mathcal{S}, \overline{a}) \models \varphi(\overline{x})$.*

One may consider the delay between two consecutive solutions as an important point in the complexity of enumeration problems. In [JYP88] several complexity measures for enumeration have been defined. One of the most interesting is that of *polynomial delay* algorithm. An algorithm $\mathcal{A}$ is said to run within a *polynomial delay* if there is no more than a (fixed) polynomial delay between two consecutive solutions it outputs (and no more than a polynomial delay to output the first solution and between the last solution and the end of the algorithm). *Polynomial delay* is often considered as the right notion of feasability for enumeration problems.
In this paper, we introduce a much stronger complexity measure that forces *constant delay* between outputs.

Definition 5 *An enumeration problem* ENUM·A *is constant delay with linear precomputation, which is written* ENUM·$A \in$ CONSTANT-DELAY$_{lin}$*, if there exists a RAM algorithm* $\mathcal{A}$ *which, for any input* x*, enumerates all the elements of the set* ENUM$\cdot A(x)$ *with a constant delay, i.e., that satisfies the following properties.*

1. $\mathcal{A}$ *uses linear input space, i.e., space* $O(|x|)$
2. $\mathcal{A}$ *can be decomposed into the two following successive steps*
   (a) PRECOMP($\mathcal{A}$) *which runs some precomputations in time* $O(|x|)$*, and*
   (b) ENUM($\mathcal{A}$) *which outputs all solutions within a delay bounded by some constant* DELAY($\mathcal{A}$)*. This delay applies between two consecutive solutions and after the last one.*

Allowing polynomial time precomputations (and polynomial space) instead of linear time, one may define a larger class called CONSTANT-DELAY$_{poly}$.

Remark 2 *As proved for the linear time class* DLIN *(see [GS02]), it can be shown that the complexity enumeration class* CONSTANT-DELAY$_{lin}$ *is robust, i.e., is not modified if the set of allowed operations and statements of the RAMs is changed in many ways. This is because linear time (and linear space) precomputations give the ability to precompute the tables of new allowed operations.*

The following result is immediate, it evaluates the total time cost of any constant delay algorithm.

Lemma 1 *Let* ENUM·A *be an enumeration problem belonging to* CONSTANT-DELAY$_{lin}$ *then, for any input* x*, the set* ENUM$\cdot A(x)$ *can be computed in* $O(|x| + |\text{ENUM}\cdot A(x)|)$ *total time, i.e., in time linear in the size of* $|Input| + |Output|$*, and linear input space* $O(|x|)$*.*

Remark 3 *In the query problem we consider, the size of* φ *is considered as a parameter. Then,* $|x| = |\mathcal{S}|$ *and the constant delay depends on* $|\varphi|$ *only.*

The two lemmas below give basic properties of constant delay computations.

Lemma 2 *An enumeration problem* ENUM$\cdot A$ *computable in linear time* $O(|x|)$ *for any input* x *belongs to* CONSTANT-DELAY$_{lin}$*.*

Proof. For any input x, one only has to compute the set ENUM$\cdot A(x)$, to sort it and to eliminate the possible multiple occurrences of solutions. These steps can be viewed as the precomputation part of the algorithm running in time $O(|x|)$. Then, one has to enumerate one by one the solutions of the sorted list. This is obviously a constant delay process. □

Lemma 3 *Let* ENUM$\cdot A$ *and* ENUM$\cdot B$ *be two disjoint enumeration problems, i.e., such that, for any input* x*,* ENUM·$A(x) \cap$ ENUM·$B(x) = \emptyset$*. Let* ENUM·$(A \cup B)$ *be the union of this two enumeration problems defined by, for any* x*:*

$$\text{ENUM}\cdot(A \cup B)(x) = \{y \ : \ A(x,y) \textit{ or } B(x,y) \textit{ holds } \}.$$

If ENUM$\cdot A$ *and* ENUM$\cdot B$ *belong to* CONSTANT-DELAY$_{lin}$ *then, problem* ENUM$\cdot A \cup B$ *also belongs to* CONSTANT-DELAY$_{lin}$*.*

Proof. Due to the disjointness of the two solutions sets for any input, the proof is evident. Given $\mathcal{A}$ and $\mathcal{B}$ the algorithms for problems ENUM$\cdot A$ and ENUM$\cdot B$, the following algorithm correctly computes for the problem ENUM$\cdot A \cup B$.

Algorithm 1 Constant delay algorithm for ENUM$\cdot A \cup B$

1: **Input:** x
2: PRECOMP($\mathcal{A}$); PRECOMP($\mathcal{B}$)
3: ENUM($\mathcal{A}$); ENUM($\mathcal{B}$)

Obviously, the delay is bounded by the maximum of DELAY($\mathcal{A}$) and DELAY($\mathcal{B}$). □

Remark 4 *Note that the disjointness condition in the Lemma above is not always necessary. In case there exist a total ordering $\leq$ and constant delay enumeration algorithms for* ENUM$\cdot A$ *and* ENUM$\cdot B$ *that enumerate solutions with respect to this unique ordering $\leq$ then, it is easily seen that* ENUM$\cdot A \cup B$ *belongs also to* CONSTANT-DELAY$_{lin}$ *even if the problems are not disjoints.*

2 First-order queries on bijective structures

2.1 Quantifier elimination on bijective structures

The key result of this paper consists of a quantifier elimination method for $\mathbf{FO_{Bij}}$ formulas.

Theorem 4 (quantifier elimination for $\mathbf{FO_{Bij}}$) *Each bijective first-order formula is equivalent to a Boolean combination of bijective atomic formulas. More precisely, let $\varphi(\overline{t}) \in \mathbf{FO_{Bij}}$ with free variables $\overline{t}$ then, there exists a Boolean combination of bijective atomic formulas $\varphi'(\overline{t})$ over the same free variables $\overline{t}$ equivalent to $\varphi(\overline{t})$.*
In the special case where φ is closed (i.e., without free variable) then, φ is equivalent to a Boolean combination of cardinality statements.

Proof. As $\forall x \varphi \equiv \neg(\exists x \neg \varphi)$, we only have to consider elimination of existentially quantified variables. W.l.o.g., we consider formulas in disjunctive normal form and, as existential quantifier commutes with disjunction we may consider the case of the elimination of a single existentially quantified variable y in a formula of the form:

$$\varphi(\overline{x}) \equiv \exists y \ (\alpha_1 \wedge \ldots \wedge \alpha_r) \tag{1}$$

where each α_i is a bijective literal among variables $\overline{x}$ and y. Literals depending on $\overline{x}$ only and cardinality statements need not be considered since they do not involve y, so $\varphi(\overline{x})$ may be supposed of the following form:

$$\varphi(\overline{x}) \equiv \exists y\ [\psi(y) \wedge y =_{\epsilon_1} \tau_1(x_{i_1}) \wedge \ldots \wedge y =_{\epsilon_k} \tau_k(x_{i_k})] \tag{2}$$

where each $y =_{\epsilon_j} \tau_j(x_{i_j})$ with $\epsilon_j = \pm 1$ is $y = \tau_j(x_{i_j})$ if $\epsilon_j = 1$ or $y \neq \tau_j(x_{i_j})$ if $\epsilon_j = -1$. To eliminate quantified variable y two cases may happen.
Suppose first there is at least one index j such that $\epsilon_j = 1$. In this case, the equality $y = \tau_j(x_{i_j})$ is used to replace each occurrence of y in the formula by the term $\tau_j(x_{i_j})$. The process results in a new formula $\varphi'(\overline{x})$ without variable y.
The second possibility leads to a more complicated replacement scheme. Suppose that for every j, $\epsilon_j = -1$. Then,

$$\varphi(\overline{x}) \equiv \exists y\ [\psi(y) \wedge \bigwedge_{j \leq k} y \neq \tau_j(x_j)] \tag{3}$$

(For simplicity of notations but w.l.o.g. we have supposed that $i_j = j$ for $j = 1, \ldots, k$). The basic idea is now the following : suppose $h \leq k$ is the number of distinct values among the k terms $\tau_j(x_j)$ such that $\psi(\tau_j(x_j))$ is true; then, formula $\varphi(\overline{x})$ is true if and only if the number of y such that $\psi(y)$ holds is strictly greater than h (i.e., $\exists_y^{h+1}\psi(y)$ is true). Introducing (new) cardinality statements in the formula, $\varphi(\overline{x})$ can be equivalently rephrased as the following Boolean combination of bijective atomic formulas:

$$\varphi(\overline{x}) \equiv \bigvee_{h=0}^{k} \bigvee_{P \subseteq [k], Q \subseteq P, |Q|=h} \left[\bigwedge_{j \in Q} \psi(\tau_j(x_j)) \wedge \bigwedge_{i \in P} \bigvee_{j \in Q} \tau_i(x_i) = \tau_j(x_j) \wedge \bigwedge_{j \in [k] \setminus P} \neg\psi(\tau_j(x_j)) \wedge \exists_y^{h+1}\psi(y) \right] \tag{4}$$

where $[k] = \{1, \ldots, k\}$.
More generally, starting from a prenex bijective first-order formula $\varphi(\overline{t})$ with free variables $\overline{t}$, one eliminates all quantified variables from the innermost to the outermost one. This will result in an equivalent Boolean combination of bijective *atomic* formulas over $\overline{t}$. In the case where φ is without free variable (i.e., $\overline{t}$ is empty), it is easily seen that the elimination process results in a Boolean combination of cardinality statements (note that, of course, $\exists x \varphi(x) \equiv \exists_x^1 \varphi(x)$). □

One interesting consequence of Theorem 4 is the following result.

Corollary 5 (Seese [See96]) *The problem* $\mathrm{MC}(\mathbf{FO_{Bij}})$ *is decidable in time* $O_\varphi(|\mathcal{S}|)$.

Proof. From Theorem 4, we know that there exists a Boolean combination of cardinality statements over the same signature σ equivalent to Φ. Given a formula $\exists_x^k \Psi(x)$ one can test whether a given σ-structure $\mathcal{S}$ satisfies $\mathcal{S} \models \exists_x^k \Psi(x)$ in time $O_\Psi(|\mathcal{S}|)$: it suffices to enumerate all the elements a of the domain, test whether $(\mathcal{S}, a) \models \Psi(x)$ in constant time and count those for which the answer is positive. If this number is greater than

or equal to k then $\exists_x^k \Psi(x)$ is true in $\mathcal{S}$. The final answer for Φ is given by the boolean combination of the answers for each cardinality statement. □

2.1.1 Considerations on an efficient implementation of the algorithm

Compared to the method of [See96], the proofs given in this paper are constructive and easily implementable. But, due to the case of Formula 3 in Theorem 4 which leads to the equivalent Formula 4 the whole process is in $O_\varphi(|\mathcal{S}|) = O(f(|\varphi|).|\mathcal{S}|)$ for some function f that may be a tower of exponentials. It can be shown that it heavily depends on the number of variables and of quantifier alternations of the formula. However, the size of the function f can be substantially reduced in case there are few quantifier alternations. In what follows, we revisit the method of the proof of Theorem 4 to prove a slightly different result in a specific case. We focus on formulas with existentially quantified variables only and show that the model-checking problem for such formulas can be efficiently evaluated. A $\mathbf{FO_{Bij}}$ formula is in $\Sigma_1 - \mathbf{FO_{Bij}}$ if it is of the form:

$$\exists \overline{y}\ \varphi$$

where φ is quantifier-free and in disjunctive normal form (DNF).

Corollary 6 *The model-checking problem for $\Sigma_1 - \mathbf{FO_{Bij}}$ formulas can be evaluated in time $O(|\varphi|^d.|\mathcal{S}|)$ where d is the number of distinct variables of φ.*

Proof. The result obviously holds for $d = 1$. So, assume $d > 1$. For the same reason as in Theorem 4, we may consider any formula of the form:

$$\varphi(\overline{x}) \equiv \exists y\ (\alpha_1 \wedge \ldots \wedge \alpha_r) \tag{5}$$

where each α_i is a bijective literal [1] with variables among $\overline{x}$ and y. For sake of completeness here, we consider also terms not containing y. Then, $\varphi(\overline{x})$ is of the form:

$$\varphi(\overline{x}) \equiv \exists y\ [\psi(y) \wedge y =_{\epsilon_1} \tau_1(x_{i_1}) \wedge \ldots \wedge y =_{\epsilon_k} \tau_k(x_{i_k}) \wedge \gamma(\overline{x})] \tag{6}$$

with the same notation ϵ_j as in the proof of Theorem 4 and $\gamma(\overline{x})$ involves variables of $\overline{x}$ only. Again, if $\epsilon_j = 1$, for some j, then all the occurences of y are replaced by $\tau_j(x_{i_j})$ and $\varphi(\overline{x})$ is equivalent to a conjunction of literals without variable y.
Suppose now that $\epsilon_j = -1$ for all $j \leq k$. Let $A = \{a \in D : (\mathcal{S}, a) \models \psi(y)\}$. Since $\psi(y)$ is quantifier-free, A can be computed in time $O(|\psi|.|\mathcal{S}|)$. Two cases need to be considered now. If $|A| > k$, since there are at most k different values $\tau_j(x_j)$ for $j = 1, \ldots, k$, then the conjunction $\exists y[\psi(y) \wedge y \neq \tau_1(x_{i_1}) \wedge \ldots \wedge y \neq \tau_k(x_{i_k})]$ is always true and $\varphi(\overline{x})$ is simply equivalent to $\gamma(\overline{x})$. If $|A| \leq k$ let $A = \{a_1, \ldots, a_h\}$, with $h \leq k$. Formula $\varphi(\overline{x})$ is replaced by the equivalent formula below over the richer signature $\sigma \cup \{a_1, \ldots, a_h\}$:

$$\bigvee_{i \leq h} \Big(\bigwedge_{j \leq k} a_i \neq \tau_j(x_{i_j}) \wedge \gamma(\overline{x}) \Big)$$

[1] In this proof, bijective literals do not involve cardinality statements

In all cases, the formula obtained is also in DNF. Time $O(|\varphi|.|\mathcal{S}|)$ is needed to eliminate variable y and the new formula is of size bounded by $O(k.|\varphi|)$, i.e., less than $O(|\varphi|^2)$. Elimination of all the d existentially quantified variables except the last one can be pursued from this new formula (without need for a normalisation). In the worst case (where all literals are of the form $x_i \neq \tau_1(x_j)$), the process will result in a disjunction of less than $|\varphi|^{d-1}$ conjunctions of at most $|\varphi|$ literals. □

2.2 Constant delay algorithm for first-order queries on bijective structures

We are now ready to state the main result of this section.

Theorem 7 *The problem* QUERY($\mathbf{FO_{Bij}}$) $\in$ CONSTANT-DELAY$_{lin}$. *In particular, from Lemma 1, it can be computed in time* $O_\varphi(|\mathcal{S}| + |\varphi(\mathcal{S})|)$ *and space* $O_\varphi(|\mathcal{S}|)$.

Definition 6 *A bijective literal is a bijective atomic formula or its negation.*

Before proving Theorem 7, we establish the following lemma.

Lemma 8 *Let S be a bijective structure and Ψ be a conjunction of bijective literals. Computing query $\mathcal{S} \mapsto \Psi(\mathcal{S})$ can be done in* CONSTANT-DELAY$_{lin}$.

Proof. The result is proved by induction on k the number of free variables of $\Psi(\overline{x})$ where $\overline{x} = (x_1, \ldots, x_k)$. We even assume that Ψ makes use of explicit constants from domain D of $\mathcal{S}$.
For the case $k = 1$, it is evident that the one variable query $Q = \{a \in D : (\mathcal{S}, a) \models \Psi(x)\}$ can be evaluated in time $O_\Psi(|D|) = O_\Psi(|\mathcal{S}|)$ and hence, by Lemma 2, is in CONSTANT-DELAY$_{lin}$.
The result is supposed to be true for k ($k \geq 1$) and proved now for $k+1$. Let's consider the query:

$$Q = \{(\overline{a}, b) \in D^{k+1} : \mathcal{S} \models \Psi(\overline{x}, y)\}$$

where the conjunction of bijective literals Ψ is over variables $\overline{x} = (x_1, \ldots, x_k)$ and y. As for Theorem 4, two cases need to be distinguished.

1. Ψ contains at least one literal of the form $\tau_1(y) = \tau_2(x_{i_0})$, $1 \leq i_0 \leq k$, that can also be rephrased as $y = \tau(x_{i_0})$,
2. Ψ does not contain such a literal.

In the first case, Ψ can rewritten as:

$$\Psi(\overline{x}, y) = \Psi_0(\overline{x}, y) \wedge y = \tau(x_{i_0}).$$

Query Q is then equivalent to:

$$Q = \{(\overline{a}, \tau(a_{i_0})) \in D^{k+1} : (\mathcal{S}, \overline{a}) \models \Psi_0(\overline{x}, \tau(x_{i_0}))\},$$

which is essentially the following k variable query Q':

$$Q' = \{\overline{a} \in D^k : (\mathcal{S}, \overline{a}) \models \Psi_0(\overline{x}, \tau(x_{i_0})\}.$$

To be precise, $Q = \{(\overline{a}, \tau(a_{i_0})) : \overline{a} \in Q'\}$. By the induction hypothesis, query Q' can be computed by some algorithm $\mathcal{A}'$ in constant delay. This provides the following constant delay procedure for query Q.

Algorithm 2 Evaluating query Q

1: **Input:** $\mathcal{S}, \Psi$
2: PRECOMP($\mathcal{A}'$)
3: Apply ENUM($\mathcal{A}'$) and for each enumerated tuple $\overline{a}$, output $(\overline{a}, \tau(a_{i_0}))$ instead

Case 2 is a little more complicated. Formula Ψ can be put under the following form:

$$\Psi \equiv \Psi_1(\overline{x}) \wedge \Psi_2(y) \wedge \bigwedge_{1 \leq i \leq r} y \neq \tau_i(x_{j_i})$$

with $1 \leq j_i \leq k$ for $1 \leq i \leq r$. By induction hypothesis, the k variable query:

$$Q_1 = \{\overline{a} \in D^k : (\mathcal{S}, \overline{a}) \models \Psi_1(\overline{x})\}$$

can be computed by an algorithm $\mathcal{A}_1$ on input $\mathcal{S}$ with constant delay. For similar reason, the k variable query Q_b over structure $(\mathcal{S}, b)$ defined by:

$$Q_b = \{\overline{a} \in D^k : (\mathcal{S}, \overline{a}, b) \models \Psi(\overline{x}, y)\}\}$$

can be enumerated by an algorithm using constant delay. Let now Q_2 be:

$$Q_2 = \{b \in D : (\mathcal{S}, b) \models \Psi_2(y)\}.$$

If $|Q_2| \leq r$ then, by Lemma 3, there exists an algorithm $\mathcal{A}_0$ which enumerates the disjoint union $\cup_{b \in Q_2} Q_b \times \{b\}$ with constant delay. Note that $\cup_{b \in Q_2} Q_b \times \{b\} = Q$. From what has been said Algorithm 3 below correctly computes query Q.

Up to step 5 of the algorithm, all can be done in linear time.
It remains to show that, in the case where $|Q_2| \geq r+1$, the delay between two successive solutions is bounded by some constant. Since $|Q_2| \geq r+1$ and the number of $b \in Q_2$ that verify $(\mathcal{S}, \overline{a}, b) \not\models \bigvee_{1 \leq i \leq r} y = \tau_i(x_{j_i})$ is bounded by r, the algorithm outputs at least one $(\overline{a}, b)$ for each $\overline{a} \in Q_1$. More precisely, it outputs $|Q_2| - r$ such tuples. For the same reasons, the maximal delay between two successive outputs is then bounded by $2r$. The same arguments apply for the delay between the last solution and the end of the algorithm. Then, computing Q can be done in constant delay. □

Proof of Theorem 7. Let $\mathcal{S}$ and $\varphi(\overline{x})$ be instances of the QUERY($\mathbf{FO_{Bij}}$) problem. From Theorem 4, one can transform $\varphi(\overline{x})$ into the following equivalent formula in disjunctive normal form:

$$\varphi(\overline{x}) \equiv \Psi_1(\overline{x}) \vee \ldots \vee \Psi_q(\overline{x})$$

where each Ψ_i is a conjunction of bijective literals and for all i, j, $1 \leq i < j \leq q$ and all bijective structures $\mathcal{S}$, $\Psi_i(\mathcal{S}) \cap \Psi_j(\mathcal{S}) = \emptyset$. The Theorem immediately follows from Lemma 3 since the enumeration problem of each query $\mathcal{S} \mapsto \Psi_i(\mathcal{S})$, $1 \leq i \leq q$, belongs to CONSTANT-DELAY$_{lin}$ by Lemma 8. □

Algorithm 3 Evaluating query Q

Input: $\mathcal{S}, \Psi$
Compute Q_2 and $|Q_2|$
if $|Q_2| \leq r$ **then** run $\mathcal{A}_0$
else
    PRECOMP($\mathcal{A}_1$)
    for $\overline{a} \in$ ENUM($\mathcal{A}_1$) **do**
        for $b \in Q_2$ **do**
            if $(\mathcal{S}, \overline{a}, b) \not\models \bigvee_{1 \leq i \leq r} y = \tau_i(x_{j_i})$ **then** Output $(\overline{a}, b)$
            end if
        end for
    end for
end if

3 Relational structures of bounded degree

3.1 Two equivalent definitions

Let $\rho = \{R_1, \ldots, R_q\}$ be a relational signature, i.e., a signature made of relational symbols R_i each of arity a_i. Recall that $arity(\rho) = max_{1 \leq i \leq q}(a_i) = m$.
Let $\mathcal{S} = \langle D; R_1, \ldots, R_q \rangle$ be a ρ-structure. For each $i \leq q$, $R_i \subseteq D^{a_i}$. The *degree* of an element x in $\mathcal{S}$ is defined as follows:

$$degree_{\mathcal{S}}(x) = \sum_{1 \leq i \leq q} \sum_{1 \leq j \leq a_i} \sharp\{(y_1, \ldots, y_{a_i}) \in D^{a_i} : \exists j \leq a_i \text{ s.t. } x = y_j \text{ and } \mathcal{S} \models R_i(y_1, \ldots, y_{a_i})\}.$$

Intuitively, $degree_{\mathcal{S}}(x)$ is the total number of tuples of relations R_i to which x belongs to. One defines the degree of a structure as $degree(\mathcal{S}) = max_{x \in D}(degree_{\mathcal{S}}(x))$.

Remark 5 *In [See96] a different definition of the degree of a structure is given. It counts, for each x, the number of distinct elements $y \neq x$ adjacent to x, i.e., that appear in some tuple with x. More precisely,*

$$degree^1_{\mathcal{S}}(x) = \sharp\{y : y \neq x \text{ and } \exists i \leq q, \overline{t} \in D^{a_i}, \text{ s.t. } \mathcal{S} \models R_i(\overline{t}) \text{ and } x, y \in \overline{t}\},$$

and $degree^1(\mathcal{S}) = max_{x \in D}(degree^1_{\mathcal{S}}(x))$.
Since each tuple containing x contains at most $m-1$ elements different from x, it is easily seen that:

$$degree^1(\mathcal{S}) \leq (m-1).degree(\mathcal{S}) \text{ where } m = arity(\rho).$$

.
Conversely, for each x, if there exist at most d elements $y \in D$ adjacent to x then, the number of distinct tuples involving x and y is bounded by $q.m.d^{m-1}$. Hence,

$$degree(\mathcal{S}) \leq q.m.(degree^1(\mathcal{S}))^{m-1}.$$

So, the two measures yield the same notion of bounded degree structure.

We are interested in the complexity of the following query problem for bounded degree structures (which is clearly independent of either measure of degree we choose).

QUERY($\mathbf{FO_{Deg}}$)
Input: an integer d, a relational signature ρ, a ρ-structure $\mathcal{S}$ with $degree(\mathcal{S}) \leq d$ and a first-order ρ-formula $\varphi(\overline{x})$ with k free variables $\overline{x} = (x_1, \ldots, x_k)$
Parameter: d, φ
Output: $\varphi(\mathcal{S}) = \{\overline{a} \in D^k : (\mathcal{S}, \overline{a}) \models \varphi(\overline{x})\}$.

3.2 Interpreting a structure of bounded degree into a bijective structure

In this section, we present a natural reduction from QUERY($\mathbf{FO_{Deg}}$) to QUERY($\mathbf{FO_{Bij}}$) which is obtained by interpreting any structure of bounded degree into a bijective one. Let $\mathcal{S} = \langle D; R_1, \ldots, R_q \rangle$ be a ρ-structure of domain D, of arity $m = max_{1 \leq i \leq q} arity(R_i)$ and of degree bounded by some constant d. One associates to $\mathcal{S}$ a bijective σ-structure $\mathcal{S}' = \langle D'; D, T_1, \ldots, T_q, g, f_1, \ldots, f_m \rangle$ of domain D' where $D, T_1, \ldots, T_q$ are pairwise disjoints unary relations (i.e. subsets of D') and $g, f_1, \ldots, f_m$ are permutations of D'. Structure $\mathcal{S}'$ is precisely defined as follows:

- D corresponds to the domain of $\mathcal{S}$.
- T_i $(1 \leq i \leq q)$ is a set of elements each representing a tuple of R_i (hence, $card(T_i) = card(R_i)$).

  The new domain D' is the disjoint union: $D \cup (D \times \{1, \ldots, d\}) \cup T_1 \cup \ldots \cup T_q$. Let us use the following convenient abbreviations: $U = D \cup (D \times \{1, \ldots, d\})$ and $T = \bigcup_{1 \leq i \leq q} T_i$.
- g creates a cycle that relates d copies of each element x of the domain. More precisely, for each $x \in D$, it holds $g(x) = (x, 1)$, $g((x, i)) = (x, i+1)$ for $1 \leq i < d$, and $g((x, d)) = x$. We also set $g(x) = x$ for all other x $(x \in T)$.
- Each f_i is an involutive permutation and essentially represents a projection of T into D as follows. Let $R_i(x_1, \ldots, x_k)$ be true in $\mathcal{S}$ for some relation R_i of arity $k \leq m$ and some k-tuple $(x_1, \ldots, x_k) \in D^k$. Suppose $R_i(x_1, \ldots, x_k)$ is represented by element $t \in T_i$, then, for each $j \leq k$, set $f_j(t) = (x_j, h)$ and set the reciprocal $f((x_j, h)) = t$ if $R(x_1, \ldots, x_k)$ is the h^{th} tuple in which x_j appears (with $h \leq d$). The construction is completed by loops $f_j(x) = x$ for all other $x \in D'$.

Figure 1 details the reduction on an example.
It is clear that, by construction, $\mathcal{S}'$ is a bijective structure and that we have the following interpretation Lemma.

Lemma 9 *Let θ_i be the σ-formula below associated to any symbol $R_i \in \rho$ of arity k:*

$$\theta_i(x_1, \ldots, x_k) \equiv \exists t (T_i(t) \wedge \bigwedge_{1 \leq j \leq k} \bigvee_{1 \leq h \leq d} f_j(t) = g^h(x_j)).$$

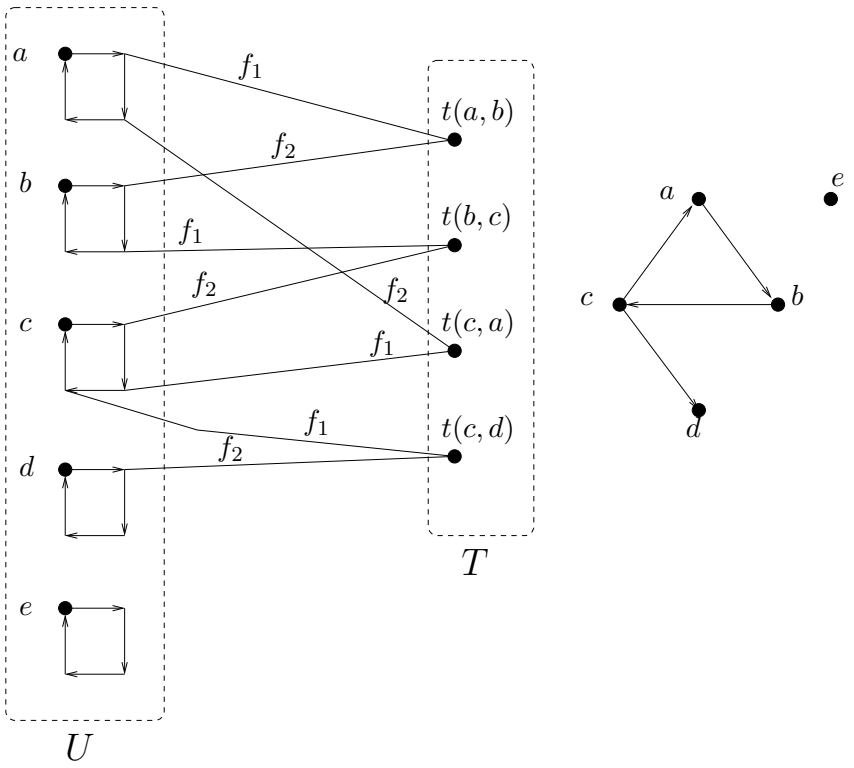

Figure 1: Our reduction on an example: the original structure (digraph) of degree 3 is on the right side of the picture

Then, for all $(a_1 \ldots, a_k) \in D^k$*:*

$$(\mathcal{S}, a_1, \ldots, a_k) \models R_i(x_1, \ldots, x_k) \iff (\mathcal{S}', a_1, \ldots, a_k) \models \theta_i(x_1, \ldots, x_k).$$

To each first-order ρ-formula $\varphi(x_1, \ldots, x_p)$, one associates the σ-formula $\varphi''(x_1, \ldots, x_p)$ obtained by replacing each quantification $\exists v$ (resp. $\forall v$) by the relativized quantification $(\exists v D(v))$ (resp. $(\forall v D(v))$) (that can be written respectively as $\exists v(D(v) \wedge ...)$ and $\forall v(D(v) \rightarrow ...)$) and by replacing each subformula $R_i(x_1, \ldots, x_k)$ by $\theta_i(x_1, \ldots, x_k)$.
The following proposition and lemma express that our reduction is correct and linear in $|\mathcal{S}|$. Because of Lemma 9, Proposition 10 can be easily proved by induction on formula φ.

Proposition 10 (interpretation of $\mathcal{S}$ into $\mathcal{S}'$) *For all* $(x_1 \ldots, x_p) \in D^p$*:*

$$(\mathcal{S}, a_1, \ldots, a_p) \models \varphi(x_1, \ldots, x_p) \iff (\mathcal{S}', a_1, \ldots, a_p) \models \varphi''(x_1, \ldots, x_p).$$

In other words: $\varphi(\mathcal{S}) = \varphi''(\mathcal{S}') \cap D^p$*. Then, setting* $\varphi'(x_1, \ldots, x_p) \equiv \varphi''(x_1, \ldots, x_p) \wedge \bigwedge_{i \leq p} D(x_i)$*, it holds:* $\varphi(\mathcal{S}) = \varphi'(\mathcal{S}')$

Lemma 11 *Computing* $\mathcal{S}'$ *from* $\mathcal{S}$ *can be done in linear time* $O_{\rho,d}(|\mathcal{S}|)$*.*

Proof. As computing $\mathcal{S}'$ from $\mathcal{S}$ is easy, one has only to compare the size of the two structures. The size of $\mathcal{S}$ is:

$$|\mathcal{S}| = \Theta(|D| + \sum_{i=1}^{q} card(R_i).arity(R_i)) = \Theta_\rho(|D| + \sum_{i=1}^{q} card(R_i)).$$

For $\mathcal{S}'$, by construction, it holds that:

$$|D'| = (d+1).|D| + \sum_{i=1}^{q} card(R_i) = \Theta_{d,\rho}(|\mathcal{S}|).$$

Hence, $|\mathcal{S}'| = \Theta(m|D'|) = \Theta_{d,\rho}(|\mathcal{S}|)$. □

We are now ready to state and prove the main result of this section.

Theorem 12 QUERY($\mathbf{FO_{Deg}}$) *belongs to* CONSTANT-DELAY$_{lin}$.

Proof. Let $\mathcal{A}$ be a constant delay algorithm that computes queries of QUERY($\mathbf{FO_{Bij}}$). By using Proposition 10, the algorithm below correctly evaluates queries in QUERY($\mathbf{FO_{Deg}}$).

Algorithm 4 Evaluating QUERY($\mathbf{FO_{Deg}}$)

1: **Input:** $\mathcal{S}, d, \varphi$
2: Compute the σ-formula $\varphi'(\overline{x})$ associated to φ (and d)
3: Compute the bijective σ-structure $\mathcal{S}'$ associated to $\mathcal{S}$ (and d)
4: Run $\mathcal{A}$ on input $\mathcal{S}'$, φ'

The cost of instruction 2 is $O_{\varphi,d}(1)$, that of instruction 3 is $O_{\varphi,d}(|\mathcal{S}|)$ (by Lemma 11) and the precomputation part of algorithm $\mathcal{A}$ (included in instruction 4) is $O_{\varphi'}(|\mathcal{S}'|)$ (hence $O_{\varphi,d}(|\mathcal{S}|)$) by Theorem 7. These steps form a precomputation phase of time complexity $O_{\varphi,d}(|\mathcal{S}|)$. Finally, the effective enumeration of $\varphi(\mathcal{S}) = \varphi'(\mathcal{S}')$ is handled on $\mathcal{S}', \varphi'$ by $\mathcal{A}$ and is performed with constant delay. □

3.3 Complexity of subgraphs problems

In this part, we present a simple application of our result to a well-known graph problem. Given two graphs $G = \langle V; E\rangle$ and $H = \langle V_H; E_H\rangle$, H is said to be a *subgraph* (resp. *induced subgraph*) of G if there is a one-to-one function g from V_H to V such that, for all $u, v \in V_H$, $E(g(u), g(v))$ holds if (resp. if and only if) $E_H(u, v)$ holds.

GENERATE SUBGRAPH (resp. GENERATE INDUCED SUBGRAPH)
Input: any graph H and a graph G of degree bounded by d
Parameter: $|H|, d$.
Output: All the subgraphs (resp. induced subgraphs) of G isomorphic to H.

The treewidth of a graph G is the maximal size of a node in a tree decomposition of G (see, for example, [DF99]). In [PV90] it is proved that for graphs H of treewidth at most

w, testing if a given graph H is an induced subgraph of a graph G of degree at most d can be done in time $f(|H|,d).|G|^{w+1}$. In what follows, we show that there is no reason to focus on graphs of bounded treewidth and that a better bound can be obtained for *any* graph H (provided G is of bounded degree). In the result below, we prove that not only the complexity of this decision problem is $f(|H|,d).|G|$ but that generating all the (induced) subgraphs isomorphic to H can be done with constant delay.

Corollary 13 *The problem* GENERATE SUBGRAPH *(resp.* GENERATE INDUCED SUBGRAPH*) belongs to* CONSTANT-DELAY$_{lin}$

Proof. The proof is given for the erate geinduced subgraph problem. Let $G = \langle V; E\rangle$ and $H = \langle V_H = \{h_1, \ldots, h_k\}; E_H\rangle$ ($|V_H| = k$) be the two inputs of the problem. Since G is of maximum degree d, we can partition its vertex set V into d sets $V^0, \ldots, V^d$ where each V^α is the set of vertices of degree α. This can be done in linear time $O(|G|)$. We proceed the same for graph H and obtain the sets $V_H^0, \ldots, V_H^d$. In case there exists a vertex in H of degree greater than d, it can be concluded immediately that the problem has no solution. Now, let Q be the following formula:

$$Q(x_1, \ldots, x_k) \equiv \bigwedge_{i<j\leq k} x_i \neq x_j \wedge \bigwedge_{V_H^\alpha(h_i)} V_G^\alpha(x_i) \wedge \bigwedge_{E_H(h_i,h_j)} E(x_i, x_j).$$

Formula Q simply checks that H is a subgraph of G and that each distinguished vertex x_i of G has the same degree as its associated vertex h_i in H. Note that formula Q only depends on H and d. The result follows now from Theorem 12. □

4 Conclusion

In this paper, we study the complexity of evaluating first-order queries on bounded degree structures and consider this evaluation as a dynamical process, i.e., as an enumeration problem. Our main contributions are two-fold. First, we define a simple quantifier elimination method suitable for first-order formulas which have to be evaluated against a bijective structure. Second, we define a new complexity class, called CONSTANT-DELAY$_{lin}$, for enumeration problem which can be seen as the minimal robust complexity class for this kind of problems and we prove that our query problem on bounded degree structures belong to this class.
There are several interesting directions for further researches. Among them, the two following series of questions seem worth to be studied:

- Which "natural" query problems belong to CONSTANT-DELAY$_{lin}$? More generally, which kind of combinatorial or algorithmic enumeration problems admit constant delay procedures ?

  The same questions can be asked for the larger class CONSTANT-DELAY$_{poly}$ of constant delay enumeration problems for which polynomial time (instead of linear time) precomputations are allowed.

- What are the structural properties of the class CONSTANT-DELAY$_{lin}$ or of the larger CONSTANT-DELAY$_{poly}$? Do they have complete problems ? Under which kind of reductions ? Could they be proved to be different from the classes of enumeration problems solvable with linear or polynomial delay ?

Acknowledgment. We thank Ron Fagin for a very fruitful email exchange that lead us to define complexity notions about constant delay computation.

References

[AHU74] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, 1974.

[BGKM00] E. Boros, V. Gurvich, L. Khachiyan, and K. Makino. Generating partial and multiple transversals of a hypergraph. In U. Montanari, J. D. P. Rolim, and E. Welzl, editors, *Proceedings 27th International Conference on Automata, Languages, and Programming (ICALP 2000), Geneva (Switzerland)*, volume 1853 of *Lecture Notes in Computer Science*, pages 588–599. Springer-Verlag, 2000.

[CH90] K.J. Compton and C.W. Henson. A uniform method for proving lower bounds on the computational complexity of logical theories. *Annals of pure and applied logic*, 48:pp.1–79, 1990.

[DF99] R. G. Downey and M. R. Fellows. *Parameterized complexity*. Springer-Verlag, 1999.

[EG95] T. Eiter and G. Gottlob. Identifying the minimal transversals of a hypergraph and related problems. *SIAM Journal on Computing*, 24(6):1278–1304, 1995.

[EGM03] T. Eiter, G. Gottlob, and K. Makino. New results on monotone dualization and generating hypergraph transversals. *SIAM Journal on Computing*, 32(2):514–537, 2003.

[FFG02] J. Flum, M. Frick, and M. Grohe. Query evaluation via tree decompositions. *Journal of the ACM*, 49(6):716–752, 2002.

[Gai82] H. Gaifman. On local and nonlocal properties. In J. Stern, editor, *Logic Colloquium'81*, pages 105–135. North-Holland, 1982.

[GO04] E. Grandjean and F. Olive. Graphs properties checkable in linear time in the number of vertices. *Journal of Computer and System Sciences*, 68(3):546–597, 2004.

[Gol94] L. A. Goldberg. Listing graphs that satisfy first order sentences. *Journal of Computer and System Sciences*, 49(2):408–424, 1994.

[GS02] E. Grandjean and T. Schwentick. Machine-independent characterizations and complete problems for deterministic linear time. *SIAM Journal on Computing*, 32(1):196–230, 2002.

[Han65] W. Hanf. Model-theoretic methods in the study of elementary logic. In L. Henkin J. Addison and A. Tarski, editors, *The Theory of Models*, pages 132–145. North-Holland, 1965.

[JYP88] D. S. Johnson, M. Yannakakis, and C. H. Papadimitriou. On generating all maximal independent sets. *Information Processing Letters*, 27(3):119–123, 1988.

[KSS00] D. J. Kavvadias, M. Sideri, and E. C. Stavropoulos. Generating all maximal models of a boolean expression. *Information Processing Letters*, 74(3-4):157–162, 2000.

[Lin04] S. Lindell. Monadic fixed-points are in linear time on bounded degree graphs. Unpublished manuscript, 2004.

[PV90] J. Plehn and B. Voigt. Finding minimally weighted subgraphs. In Springer, editor, *16th workshop on graph theoretic concepts in computer science*, volume 484 of *Lecture Notes in Computer Science*, pages 18–29, 1990.

[See96] D. Seese. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science*, 6(6):505–526, December 1996.